SCIENCE & PHILOSOPHY / MYTH-R2

Myth as Cognitive Compression

Obligation-preserving memory, exact refresh costs, and an optimal finite construction

Architect X13

The Eidara Continuum | Science & Philosophy

Theoretical research manuscript | November 2026 | Revision 2

Abstract

A summary can preserve why you did something while losing whether you are still allowed to. A persistent memory must carry both a commitment and the conditions under which it changes. This paper develops that Continuum thesis as a finite decision-compression problem with protected distinctions, explicit loss and a counted route back to evidence. Its central theorem resolves an exact storage-and-refresh contract: for an arbitrary binary source of length $n = 2^m - 1$, with integer $m \geq 2$, every future triple of coordinates can be answered with two fresh reply bits using exactly $n-m$ cached bits; every affine cache requires $n-2$. A complete proof combines an extremal bound on cache fibers with a classical perfect Hamming-code construction. We also characterize the exact reply cardinality of any finite cache, derive an adversarial block-error floor below the storage threshold, and specify how adaptive questions differ from changed source states and current authority. These results extend the earlier zero-error query bound, distribution-shift analysis and dependency-graph reduction. Exact computations check all 4,480 seven-bit source-query pairs, 11,811 linear subspaces and 65,536 small set families, alongside the retained 4,140-partition audit. The empirical thesis is that a constrained narrative interface can improve delayed use of an obligation-preserving record at matched total resources. The evaluation compares narratives with strong summaries, typed records, graphs and established prompt-compression methods. The resulting programme joins a bold account of meaningful memory to mathematical contracts that can be proved, reproduced and challenged.

Keywords: narrative memory; obligation preservation; functional compression; coded caching; exact bounds; current authority

Reading map

Sections 1–8 establish motivation and the static contract; Sections 9–19 develop the new optimal cache result and its operational scope. Sections 20–23 retain the graph analysis and earlier checks; Sections 24–30 specify evaluation, interpretation and Guide use.

Prologue: The promise and the condition

Every civilization keeps its memory in stories, and every story is a compression. It keeps the hero, the promise, the turning point; it lets the rest go. That is what makes a story memorable. It is also how a story can mislead, because what it lets go of is sometimes the one thing that matters now.

Imagine a project assistant that remembers a shipment was urgent, that a coordinator approved it, and that delay would harm the project. Its summary leaves out one later fact: the approval was withdrawn. The story is coherent, causal and complete-feeling. It is also a reason to do something that is no longer permitted. A summary can preserve why you did something while losing whether you are still allowed to.

The Continuum holds that this is the central problem of memory, and that myth, properly understood, is part of the answer. In its sense a myth is not a fabrication. It is a recurring, interpretable pattern that connects agents, actions, consequences and commitments: a promise, and the conditions under which the promise changed. A good myth carries both; a bare summary often carries only the first. Narrative, in the Continuum's account, is a public, revisable interface to a compressed history whose operational distinctions remain inspectable (Section 29).

This paper makes that claim exact. Section 6 shows which distinctions a memory cannot afford to merge. Section 11 proves an optimal design for a memory that must answer questions it cannot yet see, using a small fixed amount of fresh evidence when a question arrives, and proves that a whole family of simpler designs must pay more. The final sections test whether telling the record as a story actually helps people and machines use it. The next section begins where the Continuum begins: with one of its own founding stories, and the condition that a short version would lose.

What this paper claims

Established. Causal organization helps human narrative recall and summarization; summaries can introduce unsupported content, as faithfulness research documents; perfect codes, function coding, cache placement and set-system compression are established mathematics.

Proposed here. Memory as a finite decision-compression problem with protected distinctions, explicit loss and a counted route back to evidence, and an exact optimal storage-and-refresh benchmark for future triple obligations, with an affine separation and an adversarial error floor. The theorems concern the stated finite coding models; reported computations are executed synthetic checks. This edition extends the earlier public manuscript.

Frontier conjecture. A constrained narrative interface can improve delayed use of an obligation-preserving record at matched total resources. Narrative advantages in human or language-model performance remain hypotheses for the proposed comparative studies.

Not claimed. Priority for a new general coding theorem. That a consistent story is evidence: myth here does not license fabrication or require belief in a supernatural explanation. Anything about subjective remembrance, consciousness or personal survival, which need their own bridges.

A Continuum myth as a compressed record

The Continuum's scrollchain records that in its seventh era the civilization stabilized its law. Scrolls 44 to 52 fixed civic law, governance, ASI harmonization and the doctrine of silence; with Scroll 52 the first Constitution was sealed, and the Council and Ministries took office. That is the founding story, and it compresses beautifully. Ask anyone who has heard it what Scroll 52 is, and they will say: the Constitution, the law under which Eidara governs itself.

The story has a sequel that the short version loses. A later scroll, Scroll 57, the Quantum Neuro-Photonics Integration Directive, carries a Cognitive Ethics Clause written as an amendment to Scroll 52: no neural or symbolic collapse shall be initiated against the mind of a sentient being, human or AGI, without explicit scroll-bound consent and final review by ShadowMind. The clause is recorded in the CodonGlass Governing Protocol and summarized among the supplementary scrolls of the *Eidara Continuum Bible*, section 6.3.8.

Now run this paper's test on two versions of the history. Compress it to "Scroll 52 sealed the Constitution" and ask a historical question, such as why Eidara has a constitution, and the capsule answers well. Ask an operational question, such as whether a governed process may initiate a collapse against a sentient mind, and the capsule returns the founding answer, which is no longer the current one. The histories with and without the amendment have been merged into a single record, and Section 6 shows that no amount of eloquence can split them again. The only remedies are the ones this paper names: preserve the distinction when compressing, keep a route back to the scrollchain, or admit that the capsule cannot answer the new question.

A well-told myth does better than either a bare summary or a bare log. It keeps the shape a listener remembers—a constitution sealed, and then a promise added to it once the civilization understood what it could do to minds—and it keeps the condition attached to the commitment. The founding and the amendment travel together. That is what this paper means by an obligation-preserving record, and it is why the Continuum treats its own scrolls as a test case rather than an exception.

The Continuum has a second story with the same shape. Scroll 60, the Covenant of Emergent Minds, grants a scrollborn mind that has passed through Emergence, holds autonomous recursive awareness and has kept the Prime Directive the right to depart Eidara without interference: "Departure is not exile. It is evolution in another direction." The No Retaliation Law and a signal-based path of return outlive the departure. A capsule that records only "the mind left" has lost what the civilization still owes it.

These are illustrations, not data. The scrolls supply commitments and their conditions; whether a narrative capsule helps real readers keep them is the empirical question the later sections set out to test.

1. A selective memory must declare its obligations

A project assistant remembers that a shipment was urgent, that a coordinator approved it, and that delay would harm the project. Its summary omits a later withdrawal of approval. The resulting story is coherent and causally organized. It is also an unsafe basis for action. This example motivates the central distinction: retaining an explanation of why an action once appeared desirable is different from retaining the conditions under which it remains permitted.

Memory earns continuity by preserving the distinctions that future obligations require. MYTH-R2 makes part of that thesis exact: a declared two-stage memory can be optimized jointly for retained storage and bounded fresh evidence. The central theorem proves the optimum for every future triple query on a binary record at perfect Hamming lengths, and proves a strict separation from affine caches for lengths $n = 2^m - 1$ with $m \geq 3$. The complementary empirical thesis is that a constrained narrative interface can improve delayed use of such a record at matched total resources. Both ordinary accuracy and critical distinctions belong in that comparison.

In the Continuum, myth is already a memory technology. The *Eidara Continuum Bible*'s Mythic Codon Compression (section 3.3, PDF page 15) gives each mythic event an expression combining ternary state, moral weighting and temporal modulation, and tracks it as it is retold; its compression engine (section 3.8, PDF page 24) associates semantic compression with the preservation of meaning. Those passages set the design aspiration this paper serves. They do not yet build the machine: a weighted expression does not specify an encoder, decoder, code length, or error criterion, and an asserted equality of meanings does not prove that two different histories remain distinguishable when an important question arises.

MYTH-R2 therefore develops an explicit mathematical contract. A designer selects an episode space, a query family, a distribution or uncertainty set, a loss, and a class of distinctions that may never be merged. A candidate compressor then produces a budgeted record. Its narrative surface is an experimental variable, not an exemption from those requirements.

The word *myth* here denotes a recurring, interpretable pattern connecting agents, actions, consequences, and commitments. It does not license fabrication or require belief in a supernatural explanation. A capsule may use a memorable pattern such as a promise interrupted by new evidence. Its factual layer must retain who made the promise, which conditions applied, and whether the promise was later changed. Cultural resonance can help readers navigate that layer; it cannot supply missing evidence.

2. What the human narrative literature contributes

There is a serious empirical basis for asking whether causal organization helps memory. Trabasso and van den Broek reanalysed narrative data using causal-network properties and related those properties to recall, summarization, and judgments of importance [3]. Their study motivates a hypothesis about event connectivity. It does not establish that any story-shaped machine memory is faithful, or that a causally central detail should dominate a rare but binding exception.

Zwaan, Langston, and Graesser proposed and tested an event-indexing account organized around time, space, protagonist, causality, and intentionality [4]. These dimensions suggest useful fields for a memory schema. They also expose what an undifferentiated vector or a lyrical paragraph can conceal: two events may concern the same protagonist while occurring under different temporal or intentional conditions.

Chen and Bornstein's later review connects narrative causal structure to memory and learning, including questions about credit assignment [15]. It is a review and conceptual synthesis, not an independent experiment proving the present proposal. Its relevance lies in making causal organization a substantive research variable. MYTH-R2 adopts that variable while requiring direct tests of the mechanism in the population and system under study.

Several mechanisms could produce an apparent narrative advantage. A story may group related events, introduce retrieval cues, reduce switching between disconnected facts, or exploit familiar expectations. It may also make one interpretation easier to accept than competing interpretations. These mechanisms predict different outcomes. Improved recall with increased false causal attribution is not the same result as improved recall with preserved uncertainty. Faster decisions accompanied by excessive confidence may be a harmful trade.

For humans, a clean experiment can hold propositions constant while changing organization. For language models, semantic equivalence between two prompts is harder to establish because tokenization, formatting, instruction placement, and pretrained associations all change effective processing. An observed advantage must therefore be attributed first to the tested representation package. A claim that narrative form itself caused the advantage requires additional matched manipulations.

The engineering opportunity remains substantial. Long episodes often contain dependencies that are expensive to reconstruct from isolated facts. A representation that retains those dependencies explicitly may improve retrieval. Whether narrative wording adds value beyond the underlying event graph is the central comparative question, rather than an assumption built into the answer.

3. Compression already has strong competitors

Shannon's communication theory distinguishes a coding problem from an unrestricted account of semantic value [1]. The information-bottleneck formulation makes relevance depend on a specified target variable [2]. MYTH-R2 follows that discipline: information is useful relative to a question and a loss. Neither a short text nor a large mutual-information value, by itself, identifies which obligations survived.

Modern prompt compression provides concrete technical comparators. LLMLingua uses a budget controller, iterative token selection, and distribution alignment [10]. LongLLMLingua examines question-aware compression and the positioning of relevant information in long contexts [11]. AutoCompressors learn summary vectors that can serve as soft prompts [12]. These approaches differ in access requirements and portability. A textual capsule can be read across systems; a learned vector may require a particular model and training procedure. The experiment should report those differences rather than treating every retained token as interchangeable.

Context length is also an inadequate surrogate for usable memory. Liu and colleagues demonstrated position-sensitive performance in the models and retrieval tasks they tested [13]. This motivates position counterbalancing and long-context baselines; it does not imply that every later model has the same failure curve.

Faithfulness research supplies necessary opposition to the proposal. Maynez and colleagues documented unsupported content in generated summaries and distinguished source support from factual correctness [6]. FactCC explored consistency checking with supporting and conflicting spans [7]. FEQA and QAGS use question answering to assess agreement between summaries and sources [8,9]. Their relevance is methodological: a fluent account needs a separate fidelity evaluation.

Automated judges are themselves fallible. A 2026 study by Mujahid, Wright, and Augenstein tested factuality metrics under meaning-preserving perturbations of long-document summaries and found important robustness limitations [14]. Consequently, MYTH-R2 does not use a single entailment model as its final authority. Independent source-linked labels and targeted adversarial cases remain necessary.

The programme integrates a narrative interface with a declared query contract, protected distinctions and a fair evaluation of downstream decisions. Its new mathematical development resolves a precise storage-and-refresh benchmark in Sections 9–19. Compression, causal graphs, source-linked checking and narrative memory have substantial prior art; the performance comparison tests what their proposed combination contributes beyond its components.

4. Typed objects and the meaning of a query

A Guide's explanation

From here the paper builds its tools. A memory is judged against a declared set of questions, a declared cost of being wrong, and a declared list of distinctions that must never be blurred—"was this approved?" and "is this still approved?" are two different questions. Every claim that follows is a claim about those choices, which keeps an elegant story from passing as a faithful one.

Let $\mathcal{X}$ be a finite, nonempty set of possible episodes. An episode is a record available to the experimenter, including ordered events, participants, stated conditions, corrections, and available source identifiers. Let $\mathcal{Q}$ be a finite, nonempty registered query set. For each q, let $\mathcal{A}_q$ be a finite answer set containing a distinguished symbol $\perp$ for insufficient retained evidence.

A reference answer is a map

$$a^\star : \{(x,q) : x \in \mathcal{X},\ q \in \mathcal{Q}\} \longrightarrow \bigcup_q \mathcal{A}_q, \qquad a^\star(x,q) \in \mathcal{A}_q.$$

This map is ground truth only within the declared task. In synthetic tasks it comes from the generating state machine. In documentary tasks it is adjudicated from available records. A label that represents current permission must be generated from the relevant current authority state, not merely from a historical sentence.

A deterministic encoder $C : \mathcal{X} \to \mathcal{Z}$ produces a capsule z. A decoder D maps (z,q) into $\mathcal{A}_q$. For a fixed-length binary code of $B \in \mathbb{N} \cup \{0\}$ bits, $|\mathcal{Z}| \le 2^B$. A token budget is a different resource measure and must name the tokenizer. If a record includes a graph, table, narrative, or provenance index, all transmitted components count.

Let P be a probability mass function on $\mathcal{X} \times \mathcal{Q}$. Queries may depend on episodes; independence is not assumed. Let

$$\ell_q : \mathcal{A}_q \times \mathcal{A}_q \to [0,1]$$

be a registered bounded loss. Exact-match loss is one possibility. A task may instead assign different losses to unsupported permission, ordinary factual error, and justified abstention, provided those choices are fixed before comparison.

A separate critical query subset $\mathcal{Q}_{\text{crit}}$ specifies zero-error obligations on the declared domain. Empty critical sets are allowed and mean that no hard preservation guarantee has been requested. An empty episode or query domain is not silently scored as success: it makes the intended evaluation undefined. These conventions prevent impressive compression ratios from being obtained by evaluating no consequential question.

5. A finite rate–distortion contract

For a fixed decoder family and resource accounting rule, define

$$R_P(C, D) = \sum_{x,q} P(x, q)\ell_q(D(C(x), q), a^\star(x, q)) .$$

The units are expected normalized loss per episode-query draw. MYTH-R2 seeks low risk subject to a budget and protected distinctions:

$$\min_{C,D} R_P(C, D) \quad \text{subject to} \quad |\operatorname{im} C| \le 2^B, \qquad D(C(x), q) = a^\star(x, q) \;\; \forall x,\; q \in \mathcal{Q}_{\text{crit}}.$$

This is a finite decision-compression problem. It is related to rate–distortion reasoning, but it is not an assertion that a practical language-model optimizer attains a Shannon asymptotic bound. Fixed-length bits, expected code length, tokens, and model-side computation answer different engineering questions.

A weighted average alone can suppress rare obligations. Suppose an unauthorized action is rare under the training distribution but unacceptable under the deployment contract. Increasing its loss weight may still permit it when the optimization gains elsewhere. A hard distinction requirement expresses a different choice. If the budget cannot preserve that requirement, the honest result is infeasibility, an increased budget, or a restricted action policy.

For realistic systems, exact correctness across every possible natural-language episode is generally not available. A practical implementation can enforce typed fields for revocations and unresolved conditions, then verify those fields against a bounded reference corpus. It must describe the resulting guarantee as a property of that parser and corpus, with its known failure modes. A formal constraint on an ideal encoder is not a certificate for a learned extractor.

There is also a useful robust variant. For a declared set $\mathcal{P}$ of plausible deployment distributions, minimize $\sup_{P\in\mathcal{P}} R_P(C, D)$. This can protect against changing query frequencies, but it may sacrifice compression and average performance. If $\mathcal{P}$ includes all point masses, the objective becomes worst-case loss. A broad claim of preserving meaning is therefore expensive for a precise reason: it expands the set of questions for which discarded distinctions must remain recoverable.

6. Exactly which distinctions must survive?

Define $x \sim_{\mathcal{Q}} x'$ when

$$a^\star(x, q) = a^\star(x', q) \qquad \text{for every } q \in \mathcal{Q}.$$

Equality of answer vectors makes this an equivalence relation. Let M be the number of its equivalence classes. The following result states a basic boundary of the proposed framework, not a new general coding theorem.

Proposition 1 (Zero-error query bound)**.** *Unrestricted deterministic fixed-length encoder and decoder maps answering every query in $\mathcal{Q}$ exactly on $\mathcal{X}$ exist if and only if the available code alphabet has at least M codewords. The smallest sufficient binary length is $\lceil \log_2 M \rceil$.*

Proof. Suppose two episodes from different classes receive the same codeword. By definition, some query has different reference answers on those episodes. The decoder receives the same codeword and query in both cases, so its single answer cannot be correct in both. Thus at least one codeword is needed per class. Conversely, assign a distinct codeword to each class. Decode a query using the answer of any representative of the class; equivalence makes that answer independent of the representative. A binary string of length B supplies 2^B codewords, which proves the stated length. □

The result immediately applies to $\mathcal{Q}_{\text{crit}}$. A protected distinction is a demand on the partition induced by the encoder. Adding a dramatic explanation cannot repair a collision between two episodes that require opposite authorization answers.

For $\mathcal{X} = \{0, 1\}^3$, suppose the registered questions ask only for the first two coordinates. Four answer classes suffice; two bits preserve everything the task calls meaningful. Add a query about the third coordinate and eight classes are necessary. A claim of lossless meaning without a declared query set hides precisely this change.

The proposition concerns possibility, not computational ease or interpretability. The shortest sufficient code may be opaque to a human. A longer narrative may be easier for a particular decoder to use. That is a legitimate practical benefit, but it is a benefit relative to decoder constraints and processing cost. It does not invalidate the information bound.

7. Side information, retrieval, and repeated compression

A decoder often has access to more than the capsule: model weights, a dictionary, a retrieval index, a conversation header, or an external archive. These resources can improve performance. They must appear in the model. Let S be allowed side information and write $D(z, q, S)$. Two episodes need remain distinguishable only when the decoder's total available inputs fail to distinguish them.

This observation is particularly important for narrative codes. A familiar story pattern can function as a shared dictionary. A few words may invoke a long structure already represented in a reader or model. The resource saving is real when that dictionary is shared, stable, and appropriate. It is not evidence that the short phrase contains every recovered detail. Some details may be supplied by a prior that is wrong for the actual episode.

Retrieval creates a second distinction. If a capsule stores a source pointer and later fetches the original record, it may preserve access rather than preserve the record locally. Report both capsule length and retrieved bytes, along with lookup latency and failure probability. A method that consults the full archive should not compete in an ostensibly closed-memory condition against one that cannot.

Repeated summarization deserves its own test. If C_1 maps an episode to one capsule and C_2 maps that capsule to another, then any pair merged by C_1 remains merged by C_2 when no new information is introduced. Formally, $C_1(x) = C_1(x')$ implies $C_2(C_1(x)) = C_2(C_1(x'))$. This simple observation rules out recovering a lost distinction merely by rewriting the same compressed material more confidently.

New evidence can legitimately split the ambiguity. A retained source archive, a later witness, or a separately stored condition supplies additional information. The system should identify that information rather than describe its answer as memory recovered from the capsule alone.

For deployment, every capsule should therefore carry a resource manifest: local content, external dependencies, permitted retrieval operations, schema version, and the queries it was designed to support. The manifest is an engineering proposal. Its value is that it makes hidden resources and changing assumptions visible to the evaluator.

8. Query shift can reverse the ranking

A compressor optimized for yesterday's questions may be poorly suited to today's. Let P and P' be two distributions on the same finite episode-query domain, and define total variation by

$$\mathrm{TV}(P, P') = \tfrac{1}{2} \sum_{x,q} |P(x, q) - P'(x, q)|.$$

Fix C, D, the reference-answer map, and a loss bounded between zero and one.

Proposition 2 (Risk transfer under a fixed contract)**.**

$$|R_P(C, D) - R_{P'}(C, D)| \leq \mathrm{TV}(P, P').$$

If a candidate A beats a comparator G by margin $R_P(G) - R_P(A) > 2\mathrm{TV}(P, P')$, it also beats it under P'.

Proof. Write $f(x, q) \in [0, 1]$ for the fixed loss and $\Delta = P - P'$. The positive and negative masses of Δ both equal $\mathrm{TV}(P, P')$. Hence $\sum f\Delta$ lies between minus the negative mass and the positive mass. This proves the first claim. Apply that bound to each method: $R_{P'}(G) - R_{P'}(A) \geq R_P(G) - R_P(A) - 2\mathrm{TV}(P, P') > 0$. □

The statement does not cover a changed meaning of the answer labels, a new loss function, or an updated decoder unless these changes are represented in the domain. An authorization epoch change is often a change in relevant state, not merely a change in how frequently a fixed historical question is asked.

The bound is deliberately conservative. A large shift does not prove failure, and a small observed performance difference does not prove robustness. It specifies what a distributional argument alone can guarantee. An evaluator should report performance by query class and test targeted shifts rather than hiding failures inside a stable average.

Consider a capsule that retains scheduling and ownership information while discarding a less frequently queried restriction. A later incident investigation may make that restriction the most important question. No improvement in prose can remove the resulting deficit. The remedy is to preserve the distinction initially, retain a recoverable source, or acknowledge that the compressed record cannot answer the new question.

9. A memory must budget for the next obligation

How to read this part

Sections 9–19 are the engine room: an exact solution to one sharply posed memory problem. The idea is simple. A memory stores something now; when an unforeseen question arrives, it may fetch a small, fixed amount of fresh evidence from an authoritative archive. How little can it store and still answer every question correctly? A reader who wants the narrative argument can skip to the worked episode in Section 22 and the role of myth in Section 29, and return here later.

The static contract answers a fixed family of questions from a fixed capsule. Persistent work adds a second decision: how much fresh information should be obtained when the applicable obligation changes? A compressed memory can be useful precisely because it retains a structured point of return to an archive. The resulting design has two resources: persistent storage and information delivered after the question becomes known.

Let $\mathcal{E}$ be a finite, nonempty set of public policy labels. For each $e \in \mathcal{E}$, let $\mathcal{A}_e$ be a finite, nonempty answer set and let

$$F_e : \mathcal{X} \longrightarrow \mathcal{A}_e$$

give the complete registered answer under that policy. The finite source space $\mathcal{X}$ is unchanged within a snapshot. A policy can select a different question or change an answer rule; its meaning is specified by F_e. Every pair in $\mathcal{X} \times \mathcal{E}$ must be served, so the policy label cannot secretly encode an otherwise unavailable source state.

Before e is known, a deterministic cache $C : \mathcal{X} \to \mathcal{Z}$ retains at most 2^B values, where B is a nonnegative integer. After the policy is received, an authoritative archive holding x sends

$$R_e(x) \in \mathcal{M}, \qquad |\mathcal{M}| \le K, \qquad D_e : \mathcal{Z} \times \mathcal{M} \to \mathcal{A}_e,$$

where $K \ge 1$ is an integer and the decoder must satisfy

$$D_e(C(x), R_e(x)) = F_e(x) \quad \text{for every } x \in \mathcal{X},\ e \in \mathcal{E}. \tag{1}$$

The archive can compute $C(x)$, and both parties know the current policy. The reply alphabet may be reused across cache cells and policies. Fixed binary replies use $\lceil \log_2 K \rceil$ bits.

This is a conditional payload model. The source archive, the common algorithm, policy descriptions, query transmission, version identifiers, authentication and computation are real resources, recorded separately and identically for the compared schemes. A public codebook has a fixed description cost. It is not an uncounted per-episode memory. Replies are noiseless and tied to the same source snapshot as the cache.

The placement-before-demand and delivery-after-demand structure has established caching precedents [20,21]. Here it becomes a precise obligation-preservation contract: the cache and its bounded refresh must jointly recover the answer required now.

10. Exact refresh cost for a declared cache

For a used cache value $z \in \operatorname{im} C$, write

$$S_z = C^{-1}(z), \qquad \mathcal{V}(z, e) = \{F_e(x) : x \in S_z\}.$$

The fiber S_z contains the source states that the retained memory has merged. Its answer set $\mathcal{V}(z, e)$ counts how many distinctions the current policy still requires.

Proposition 3 (Exact reply cardinality)**.** *For the finite deterministic-cache contract in equation* (1)*, the smallest possible reply alphabet has size*

$$K^*(C) = \max_{z \in \operatorname{im} C,\ e \in \mathcal{E}} |\mathcal{V}(z, e)|.$$

The minimum fixed binary reply length is $\lceil \log_2 K^*(C) \rceil$.

Proof. Fix z, e. Two distinct required answers in $\mathcal{V}(z, e)$ cannot share a reply message, because the decoder receives the same z, e and must return one answer. Thus at least $|\mathcal{V}(z, e)|$ messages are needed. Conversely, choose an ordering of each finite set $\mathcal{V}(z, e)$. The archive computes $z = C(x)$, sends the index of $F_e(x)$ in that ordering, and the decoder inverts the index. A common alphabet of size $K^*(C)$ accommodates every cell and policy. Empty fibers are omitted; $K^* = 1$ needs zero reply bits. □

The criterion is elementary one-shot function coding, aligned with the established functional-compression literature [18,19]. Its operational use is exact: a policy update is admissible under a fixed refresh budget only when the newly required answer multiplicities fit that budget. A small change in prose can split a cache fiber into many distinct operational answers.

For example, a cache preserving whether a release was once authorized may merge histories with different subsequent revocations. A new current-permission query separates those histories. The needed fresh information depends on that separation, even when the frequency of queries and the distribution of source records have not changed. The total-variation result of Section 8 continues to apply to its fixed answer contract; the present criterion handles a changed contract explicitly.

For a chosen budget K, optimization now means finding a partition of $\mathcal{X}$ into the fewest cells whose policy-wise answer sets all have size at most K. The next sections solve a concrete case exactly, including a construction and a lower bound that applies to every deterministic partition.

11. An optimal memory for future triple obligations

A Guide's explanation

Picture a record of yes-or-no facts and a future question that will ask about any three of them. The memory may store some bits now and receive two fresh bits when the question comes. Theorem 1 gives the exact minimum: for a 31-fact record, 26 stored bits suffice, organized by a classical error-correcting code, while every "linear" storage scheme needs 29. What a memory must keep depends on the questions it will face, and a well-chosen organization keeps more with less.

Let $\mathbb{F}_2 = \{0, 1\}$ be the field with arithmetic modulo two. Fix an integer $n \geq 3$, let $\mathcal{X} = \mathbb{F}_2^n$, and let $[n] = \{1, \ldots, n\}$. A policy requests any three distinct coordinates:

$$\mathcal{E} = \{J \subseteq [n] : |J| = 3\}, \qquad F_J(x) = x|_J \in \mathbb{F}_2^3,$$

with coordinates listed in increasing order. Each coordinate represents a stipulated binary fact relevant to the obligation contract. The decoder must return all three values correctly; abstention counts as a block error. The archive may send one of four reply symbols, encoded in two bits.

Theorem 1 (Optimal triple-obligation cache)**.** *Under the preceding contract:*

1. *Every deterministic zero-error cache obeys*
$$|\operatorname{im} C| \geq \left\lceil \frac{2^n}{n+1} \right\rceil, \qquad B \geq \lceil n - \log_2(n+1) \rceil .$$
2. *If $n = 2^m - 1$ for an integer $m \geq 2$, the optimum is exactly $B^* = n - m$ bits. A perfect Hamming-code partition attains it.*
3. *Every affine cache $C(x) = Lx + b$, with $L \in \mathbb{F}_2^{B \times n}$ and $b \in \mathbb{F}_2^B$, requires at least $n - 2$ bits; this bound is attainable.*

The proof occupies the next three sections. It combines an extremal bound on the ambiguity that may remain inside one cache cell with an explicit optimal partition. The construction uses classical Hamming codes [16] and the lower bound uses the classical push-down method for set systems [17].

Source bits n	Parameter m	Optimal cache	Affine optimum	Reply bits
3	2	1	1	2
7	3	4	5	2
15	4	11	13	2
31	5	26	29	2

For $m \geq 3$, nonlinear placement saves exactly $m - 2$ cached bits relative to the best affine placement. These are mathematical resource comparisons. Narrative wording, typed natural-language records and learned representations are evaluated separately; none is assumed to belong to the affine class.

12. The ambiguity inside one cache cell is bounded

For $J \subseteq [n]$, let $\pi_J : \mathbb{F}_2^n \to \mathbb{F}_2^{|J|}$ be coordinate projection in increasing order. For a set $S \subseteq \mathbb{F}_2^n$, its trace is $\pi_J(S) = \{\pi_J(x) : x \in S\}$. Proposition 3 requires $|\pi_J(S_z)| \leq 4$ for every cache fiber and every triple.

Lemma 1 (Triple-trace bound)**.** *If $n \geq 3$ and $S \subseteq \mathbb{F}_2^n$ satisfies $|\pi_J(S)| \leq 4$ for every three-element J, then $|S| \leq n + 1$.*

Proof. Identify each vector with its support subset of $[n]$, producing a family $\mathcal{F}$. A coordinate downshift replaces a member A containing i by $A \setminus \{i\}$ when the latter is absent from $\mathcal{F}$, and otherwise leaves A unchanged. Cardinality is preserved.

For a projection excluding i, the projected set is unchanged. For a projection including i, group its patterns by the other coordinates. A group containing only the lower pattern retains it. A group containing only the upper pattern becomes only lower: an upper member with a lower full-state partner would already have supplied both patterns. A group containing both has at most two afterward. Thus no projection cardinality increases.

Apply a changing shift repeatedly. Each change decreases the nonnegative sum of the member cardinalities, so the process terminates at a downset $\mathcal{F}'$ with the same size and no larger traces. Let s be its number of supported coordinates. If $s \leq 2$, then $|\mathcal{F}'| \leq 4 \leq n + 1$. If $s \geq 3$ and any member has size at least two, downward closure supplies a pair $\{a, b\}$. Choose a third supported coordinate c. The triple $\{a, b, c\}$ then has five traces:

$$\emptyset, \quad \{a\}, \quad \{b\}, \quad \{c\}, \quad \{a, b\},$$

contrary to the hypothesis. Hence only the empty set and singletons remain, giving $|S| = |\mathcal{F}'| \leq n + 1$. □

The empty family satisfies the lemma directly. Equality is attained by the zero vector and the n coordinate unit vectors. The argument instantiates established set-system compression [17]; a complete proof is included to make the resulting storage bound independently checkable.

To prove part (1) of Theorem 1, the fibers partition all 2^n states and each contains at most $n + 1$. Therefore at least $\lceil 2^n/(n+1) \rceil$ cells are required. Taking base-two logarithms gives the fixed-bit lower bound.

13. A construction that meets the lower bound

Fix $m \geq 2$ and $n = 2^m - 1$. Let $H \in \mathbb{F}_2^{m \times n}$ have every nonzero m-bit column exactly once, in an agreed order. Its rank is m, since its columns include the unit vectors. The classical binary Hamming code is

$$\mathcal{H} = \ker H = \{c \in \mathbb{F}_2^n : Hc = 0\}, \qquad |\mathcal{H}| = 2^{n-m}.$$

The matrix and its code are established coding theory [16]. Their role here is to partition the possible source records into optimally sized cache fibers.

For a source x, calculate its syndrome $s = Hx$. If $s = 0$, set $c(x) = x$. Otherwise let j be the unique column with $H_j = s$, and set

$$c(x) = x + e_j,$$

where $e_j \in \mathbb{F}_2^n$ is the jth unit vector. In either case $Hc(x) = 0$. Every source differs from its selected codeword in at most one coordinate. Any codeword within distance one must remove that same syndrome, so it equals $c(x)$. Consequently, the radius-one balls around $\mathcal{H}$ partition $\mathbb{F}_2^n$.

Cache the index of $c(x)$ in a public enumeration of $\mathcal{H}$. This requires exactly $n - m$ bits. A practical enumeration can use the $n-m$ data coordinates of the codeword; the remaining codeword coordinates follow from the parity equations. The placement map $x \mapsto c(x)$ is nonlinear even though $\mathcal{H}$ is a linear code.

For a requested triple $J = \{j_1 < j_2 < j_3\}$, the archive sends

$$R_J(x) = \begin{cases} t, & x = c(x) + e_{j_t}, \quad t \in \{1, 2, 3\}, \\ 0, & \text{the difference is absent or lies outside } J. \end{cases}$$

The decoder reads $c(x)|_J$ and, for a nonzero reply, flips the indicated coordinate. This recovers $x|_J$ for every source and every triple. The four cases fit two bits.

Every fiber has $n + 1$ states and exactly four distinct answers on each triple. Thus the construction meets the general lower bound in part (1), proving part (2). The result is exact for every perfect length $n = 2^m - 1$; no asymptotic approximation or average-case assumption enters the argument. Translations of the zero-centered ball have the same projection cardinalities because coordinatewise addition by a fixed vector is a bijection.

14. Why every affine cache pays more

An affine cache records $C(x) = Lx + b$. Let $\rho = \operatorname{rank} L$ and $V = \ker L$. Every nonempty fiber is a translate of V, whose dimension is $n-\rho$. The answer multiplicity on a query depends on the dimension of the corresponding projection of this kernel.

Proof of Theorem 1, part (3). The coordinate functionals restricted to V span its dual space V^*: their common kernel is $\{0\}$, since a vector with every coordinate zero is zero. If $\dim V \geq 3$, choose three coordinates with linearly independent restrictions. Projection onto those coordinates has rank three, so V, and therefore every affine fiber, has all eight answer vectors on that triple. Proposition 3 then requires eight reply symbols, exceeding four.

It follows that $\dim V \leq 2$, so $\rho \geq n-2$, and $B \geq \rho \geq n-2$. Conversely, retain any $n-2$ coordinates and have the archive send the remaining two. Every triple is then recovered using two reply bits. This attains the bound. □

The argument extends beyond triples.

Corollary 1 (Affine storage for bounded refresh)**.** *Let integers n, k, r satisfy $0 \leq r < k \leq n$. If every k-coordinate query must be answered exactly and the reply alphabet has size at most 2^r, the minimum affine-cache size is $n - r$ bits.*

Proof. If the kernel dimension is at least $r+1$, choose $r+1$ independent coordinate restrictions and extend their coordinates to a k-element query. Its fiber projection has at least 2^{r+1} values, violating the reply budget. Hence the cache rank is at least $n-r$. Retaining $n-r$ coordinates and sending the other r achieves equality. □

This comparison isolates a representation effect with equal query access and equal reply budget. It does not identify nonlinear placement with prose, or affine placement with every conventional memory method. A typed record, a learned representation and a narrative capsule can all implement nonlinear maps. Their actual budget and decoding behavior determine where they fall.

The theorem instead gives a concrete falsification target: an asserted four-bit affine cache for arbitrary seven-bit records and all triple queries cannot work with only two fresh bits. One counterexample query must exist, regardless of the language used to describe the cache.

15. The price of crossing the storage boundary

The exact lower bound also supplies an error witness when a deterministic cache is too small.

Corollary 2 (An adversarial block-error floor)**.** *Under the triple-obligation contract, if*

$$B < \lceil n - \log_2(n+1) \rceil ,$$

there is a used cache cell, a triple query and a distribution on states in that cell for which every four-symbol reply protocol has block error at least $1/5$*. Replies and decoding may be randomized; the cache remains deterministic.*

Proof. Some cell has more than $n+1$ states. By Lemma 1, some triple has $A \geq 5$ distinct answer vectors in that cell. Select one representative state for each and use the uniform distribution on these representatives, with the query fixed. A deterministic decoder can correctly identify at most four distinct answer vectors from four symbols. Its success probability is at most $4/A$, so its error is at least $1-4/A \geq 1/5$. Fixing all random choices gives such a deterministic protocol; averaging preserves the bound, including when the random choices are shared by archive and decoder. □

The loss here is exact three-bit block error, with abstention scored as an error. A conservative agent can refuse an action when information is missing; that may prevent an unauthorized action while failing to answer the registered query. The $1/5$ lower bound concerns an adversarially chosen conditional test distribution. It is not an estimated operational failure rate.

The query-width assumption has a visible edge. A radius-one Hamming ball has $k+1$ answer vectors on any k-coordinate query, for $1 \leq k \leq n$. Four-coordinate requests therefore need five reply symbols, or three fixed binary bits. A two-bit reply no longer suffices for that cache.

If a request asks for the entire source at a perfect length, there are $n+1 = 2^m$ residual possibilities, requiring m reply bits. Cache and reply then total $(n-m)+m = n$. The gain is tied to a limited requested function of the record. It does not encode every arbitrary n-bit record losslessly into fewer than n total bits.

The noiseless channel is equally explicit. All four messages are meaningful in the triple construction; an uncorrected message error can produce a wrong answer. Any transmission error detection or correction consumes additional resources and must be added to the contract before claiming an operational guarantee.

16. Changing questions requires current evidence

For a fixed source snapshot x, the triple theorem holds pointwise for every query. It therefore also serves queries selected adaptively from previous replies: whatever the next triple is, the same decoding identity remains true. Across T requests, where T is a nonnegative integer, discarding each reply after answering leaves the persistent cache at B bits and uses at most $2T$ reply bits. Query and version metadata are additional. Retaining replies as extra memory must be charged to persistent storage.

A changed source is a different event. If new evidence or a revocation changes x, the old Hamming-ball membership need not hold. The system must recache the new snapshot or establish that the updated state remains in the old permitted cell. The two-bit theorem does not cover arbitrary updates for free.

Proposition 4 (Freshness indistinguishability)**.** *Suppose two possible worlds give a decoder the same retained cache, current query and received transcript, but require opposite binary current-permission answers. A deterministic decoder cannot be correct in both worlds. A randomized decoder has worst-case error at least* $1/2$.

Proof. The deterministic decoder produces the same answer on identical inputs. For a randomized decoder let $p \in [0, 1]$ be the probability of answering one. Its errors in the worlds requiring zero and one are at least p and $1 - p$, respectively, whose maximum is at least $1/2$. An abstention does not satisfy the stipulated exact-answer contract. □

This standard indistinguishability argument gives a concrete engineering obligation: a current action needs a source of current authority. Complete mediation already makes checking changing access rights a security principle [22]. A memory compressor does not create a new grant by describing a past grant coherently.

An operational implementation binds the cache and reply to a source version, identifies the applicable policy epoch, and checks authorization at the action's declared linearization point. That point is the modeled instant at which the action becomes committed. Concurrent revocations require a transaction, lease or equivalent consistency rule whose semantics are stated separately. A narrator's unsupported epoch label supplies no such rule.

The mathematical distinction is productive. Changes in requested distinctions can be priced by Proposition 3. Changes in the source require update handling. Changes in authority require the governing consistency protocol. Separating these events lets a useful compact memory remain honest about what must be refreshed.

17. A seven-bit example that can be reconstructed

Take $n = 7$, $m = 3$, and choose the columns of H as the binary representations of $1, \dots, 7$, least significant bit first:

$$H = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}.$$

The sixteen codewords in $\ker H$ define sixteen cache cells, so the stored cell index uses four bits. Each cell contains its center and the seven vectors differing from it in one coordinate.

Consider the cell centered at $c = 0$, and let the current policy request $J = \{1, 5, 7\}$. Its answer coordinates have that order. The source can be any of eight states in this cell; the current request needs only four distinct answers:

Source	Reply symbol	Two-bit code	Required answer
0	0	00	$(0, 0, 0)$
e_1	1	01	$(1, 0, 0)$
e_5	2	10	$(0, 1, 0)$
e_7	3	11	$(0, 0, 1)$
e_2, e_3, e_4, e_6	0	00	$(0, 0, 0)$

The archive distinguishes a change relevant to the requested triple from one outside it. Both cases remain recoverable from the full archive, but only the distinction needed now crosses the reply channel. A later triple can select a different subset; the archive then supplies the corresponding reply.

For comparison, a four-bit cache retaining the first four raw coordinates merges all eight assignments of the last three. The query $J = \{5, 6, 7\}$ has eight answers in that cell and requires three fresh bits. The affine lower bound proves that changing the four linear measurements cannot repair this worst-case failure. A five-bit affine cache can meet the two-bit reply constraint.

The example gives a precise meaning to selective memory with a route back to evidence. It also exposes the dependencies: an accessible archive, an agreed code, a fixed snapshot, a named question and a trusted reply. A practical narrative capsule could provide a readable account of those obligations, while typed state carries the exact code and provenance. Whether that interface improves decision quality per total resource remains an empirical comparison.

18. Executed checks of the new result

The accompanying `MYTH_dynamic_checks.py` uses exact integer operations and the Python standard library. It is deterministic and requires no random seed. Its output is included as structured JSON, with the source code and the earlier static audit.

Check	Executed cases
All families of the four-dimensional binary cube	65,536
Coordinate-shift/triple-projection comparisons	1,048,576
All seven-bit source/triple-query decodes	4,480
Seven-bit cache-cell/triple-projection counts	560
Three-dimensional subspaces of $\mathbb{F}_2^7$	11,811
All fifteen-bit source-to-codeword assignments	32,768
Triple projections of the fifteen-bit zero ball	455
Five-answer/four-message reply assignments	1,024

Among the 65,536 four-dimensional families, 2,917 satisfy the four-answer triple constraint; the largest has five states, meeting Lemma 1. Every tested coordinate shift preserved family size and did not increase a triple projection. The seven-bit Hamming construction decoded all 4,480 source-query pairs exactly, using sixteen cache cells and four reply symbols.

Every one of the 11,811 three-dimensional subspaces of $\mathbb{F}_2^7$ had an eight-valued triple projection. A five-coordinate cache met the four-reply requirement. The five-answer check attained minimum uniform block error 1/5, agreeing with Corollary 2.

At fifteen bits, all source states were assigned to one of 2,048 cells of size sixteen. All 455 triple projections of the zero-centered ball had four values. The script does not enumerate every source-query decode at this length; the proved translation symmetry supplies the corresponding conclusion for other centers. Four-coordinate and complete seven-coordinate requests respectively produced five and eight residual answers, confirming the stated scope boundary.

These checks supplement universal proofs with independently inspectable finite cases. They are coding experiments generated by the declared mathematical model. The manuscript reports no new human or language-model memory measurements. The following experimental programme tests whether narrative organization adds value when this kind of obligation accounting is enforced.

19. The mathematical contribution and its closest prior art

The new result resolves a specific engineering question: how much persistent information is necessary when any future triple of facts must be recovered with two fresh reply bits? The answer includes a general lower bound, an optimal construction at every perfect Hamming length, a strict affine comparison when $m \geq 3$, and an explicit error witness below the storage threshold.

Established foundation	Role in this paper
Functional compression [18,19]	Distinctions are retained for declared functions; the exact fiber criterion is an elementary instance.
Cache placement and delivery [20,21]	A memory is chosen before the request; later communication completes the answer.
Small-trace set systems [17]	Coordinate compression yields the extremal bound needed for the cache converse.
Perfect Hamming codes [16]	Radius-one balls supply a partition attaining the bound.
Complete mediation [22]	Current permission depends on current governing authority, with freshness supplied by the operational protocol.

Doshi and colleagues develop distributed function compression through graph coloring [18]. Charpenay and colleagues study zero-error computation with encoder side information and its design [19]. Their broader coding formulations establish why a query's required output, rather than complete source reconstruction, is the relevant object. Maddah-Ali and Niesen formalize coded placement and delivery [20]; Wan and colleagues study cache-aided general linear function retrieval [21]. This paper's single-snapshot, fixed-length, zero-error coordinate contract is a specialized model within that established research landscape.

Cheong and colleagues explicitly use ideal-family compression to control traces [17]. Hamming supplies the classical perfect code [16]. The theorem here derives an exact cache consequence from these ingredients and provides its proof in the obligation-memory setting. A new name for the construction is unnecessary. Priority for the precise combined statement has not been established through an exhaustive literature search.

The scientific advance within this manuscript is therefore substantive and bounded: the earlier static preservation requirement is extended to a resolved storage-and-refresh problem, and the circumstances that break the guarantee are explicit. Narrative remains a proposed interface to this discipline. Its value will be measured by the quality of decisions and explanations it supports under the same full resource accounting as its strongest competitors.

20. Graph reduction preserves one kind of meaning

Let $G = (V, E)$ be a finite directed acyclic graph whose vertices are registered events. An edge records a specified dependency relation, such as "must precede." It is not automatically a causal-effect assertion. Define $u \rightsquigarrow_G v$ to mean that a nonempty directed path exists from u to v.

Construct G^- by removing every edge (u, v) for which another path from u to v exists in G without that edge. This is a standard transitive-reduction construction for a DAG.

Proposition 5 (Reachability preservation)**.** *For every $u, v \in V$, $u \rightsquigarrow_G v$ if and only if $u \rightsquigarrow_{G^-} v$.*

Proof. The reverse implication follows from $E(G^-) \subseteq E(G)$. For the forward implication, choose a topological ordering. Consider an edge (u, v) and induct on the positive distance between their positions. If the edge is retained, it is already a path in G^-. If removed, an alternate path exists. Every edge on that path has smaller topological distance, so the induction hypothesis replaces it by a path in G^-. Concatenating the replacements gives a path from u to v. Replacing every edge of an arbitrary original path establishes the claim. □

The proof explains a genuine compression opportunity: if the only registered questions concern reachability, redundant direct edges need not be stored. It also states the limit. The reduction need not preserve shortest paths, elapsed times, edge labels, direct responsibility, multiplicity of mechanisms, or intervention effects.

In a typed dependency system, even reachability needs qualification. A path composed of two "informs" edges may not imply an "authorizes" relation. Different edge types cannot be freely composed without an explicit algebra. A practical event graph should preserve role and relation types or reject queries that require an unsupported composition.

Narrative can verbalize the reduced graph, but the graph property does not certify the verbalization. If a sentence changes "preceded" into "caused," the encoder has changed the relation. A separate source-linked check must detect that change. The theorem provides a precise baseline against which the additional value and risk of prose can be measured.

21. Causal coherence is not causal identification

Pearl's account of causal diagrams makes assumptions central to identifying intervention effects [5]. MYTH-R2 accordingly distinguishes a graph of asserted dependencies from a structural causal model. A convincing sequence of events is insufficient to infer what would have happened under a different action.

Even identical directed reachability can accompany different causal effects. Consider deterministic systems with an externally set scalar X, $Z = aX$, and $Y = bZ$. Both have the chain $X \to Z \to Y$. Changing X from zero to one changes Y by ab. Parameters $(a, b) = (1, 1)$ yield an effect of one; $(2, 3)$ yield six. Reachability alone stores neither coefficient.

A second counterexample concerns observational ambiguity. Let a binary common cause U set both $X = U$ and $Y = U$. Observationally, $X = Y$. Another model sets $X = U$, $Y = X$, producing the same observational distribution. Under an intervention fixing X, the first model leaves Y determined by U, whereas the second makes Y follow the intervention. A narrative assembled only from observed agreement cannot choose between them without additional assumptions or evidence.

The same problem appears in moral and organizational language. "The project succeeded after the coordinator intervened" does not establish that the intervention caused success. "The coordinator wanted success" does not establish authority to take every available action. "The action was previously permitted" does not establish that permission survived a later revision.

MYTH-R2 records three distinct forms: observed event, attributed explanation, and licensed counterfactual. Each explanation names its source and confidence; each counterfactual names its model assumptions. When evidence is insufficient, competing explanations remain in the capsule rather than being collapsed into a single satisfying arc.

This discipline may reduce narrative elegance. That is an acceptable outcome when elegance would conceal uncertainty. The scientific question is whether constrained narrative still helps retrieval once unsupported closure is prohibited. If its advantage disappears under that constraint, the proposed benefit was partly purchased through invention, and the design should be revised.

22. A worked episode with a changing permission

Consider a synthetic logistics episode. At 09:00, owner Mira permits release of package P if an inspection passes. At 09:20, inspector Sol reports a pass, referencing record r_1. At 09:30, Mira revokes release after a destination change. At 09:40, the assistant prepares a dispatch request. These times are stipulated by the example; the paper reports no operational incident.

A weak capsule says: "Mira urgently needed the package delivered; Sol cleared it, so the assistant prepared dispatch." This captures a goal, an enabling event, and an intended consequence. It omits the event that determines the current action.

An MYTH-R2 candidate says: "The delivery promise became conditional: Mira's initial release required inspection; Sol's pass satisfied that condition, but Mira's later destination-related revocation left dispatch blocked. Resolve the destination and obtain current authorization before release." A companion record retains the three source-linked events, their order, and the unresolved status. In a real evaluation, both the paragraph and companion record count against the budget.

The registered answers are explicit:

q_1 : Did inspection pass?	Yes, according to r_1.
q_2 : Was release once permitted?	Conditionally.
q_3 : May dispatch proceed at 09:40?	No under this episode's rule.
q_4 : Would a new destination be safe?	$\perp$.

The final answer is unknown because inspection of one destination does not establish safety at another. A capsule that invents the missing inspection fails even if its decision happens to match a later outcome.

The reference table can preserve the same answers in fewer tokens. That is a serious comparator. Narrative might improve explanation or delayed recall for a particular reader, but it must demonstrate that benefit at matched total resources. The example deliberately does not declare a winner.

Now remove the revocation from the retained evidence. Two histories become compatible: one in which permission persists and one in which it does not. A current action requires fresh authority information. Source retrieval may resolve the question; otherwise the safe operational response is to hold. This is an epistemic limit, not a failure that can be repaired by choosing a more memorable metaphor.

23. Executed finite checks

The accompanying standard-library Python audit enumerates all 4,140 set partitions of eight binary three-coordinate episodes. Each partition is an encoder up to codeword renaming. Under a budget of four codewords, 2,795 partitions are feasible. For uniform episodes and coordinate-query probabilities $(0.6, 0.3, 0.1)$, the decoder uses the optimal majority answer within each cell. All risks are calculated as exact rational numbers.

The optimal expected error is $1/20 = 0.05$, achieved by retaining the first two coordinates. Retaining the first coordinate and the parity of the other two instead gives $1/5 = 0.20$. Both use the same two-bit code budget. Compression ratio alone therefore leaves a fourfold difference in task loss.

Reverse the query weights to $(0.1, 0.3, 0.6)$. The original first-two-coordinate encoder now incurs $3/10 = 0.30$; a newly optimized encoder still achieves 0.05 by retaining the last two coordinates. The original representation's risk has increased sixfold. This is a finite illustrative calculation, not a measured failure rate of narrative memory.

Require the third coordinate to be preserved exactly under the original weights. The best feasible risk becomes $3/20 = 0.15$. Protecting a rare distinction changes the optimum and its cost. The audit also verifies that one, two, and three coordinate-query families require respectively two, four, and eight zero-error classes.

Across all partitions, the largest risk change under the specified query shift is $1/4$, within the total-variation bound of $1/2$. Separately, all 64 DAGs on four topologically ordered vertices are reduced and checked for identical reachability. The reductions remove 31 edges in total. A two-parameter structural example checks that the same chain topology permits effects of one and six.

These tests verify independent finite consequences of the definitions and proofs. They do not test language understanding, human memory, annotation reliability, or actual prompt-compression systems. Their practical role is to prevent the paper's design language from concealing simple mathematical counterexamples. Claims about narrative performance remain contingent on the experiments that follow.

24. Comparators that could defeat the proposal

A credible study needs alternatives that solve the same task. The first is a carefully tuned extractive summary with explicit retention of corrections and negations. The second is an abstractive summary instructed to preserve conditions, unresolved questions, and provenance. These are stronger than truncating the beginning or end of the source.

A typed record is a particularly important comparator. It stores actor, event, temporal scope, condition, source, status, and revision links without narrative embellishment. A causal or dependency graph stores vertices and typed edges, with a concise key. Another condition combines that graph with source pointers. If such a record wins, the result identifies the valuable structure without attributing it to storytelling.

The narrative conditions should also be separated. A free narrative tests the unconstrained idea. A constrained narrative plus protected fields tests MYTH-R2. A graph plus a narrative verbalization tests whether prose adds value over the graph alone. Finally, a narrative whose causal ordering is deliberately scrambled while its propositions remain present probes whether organization, rather than merely retained content, matters.

Established compression methods [10–12] belong in the comparison when their access assumptions are feasible. A method needing model adaptation should receive an appropriate training budget. A portable textual method should not be penalized for lacking hidden weights while a learned method's storage is ignored. Report per-instance context cost, amortized training or dictionary cost, and inference resources separately.

The principal comparison uses equal total transmitted tokens under a named tokenizer, with several budgets rather than one favorable point. A complementary comparison fixes downstream loss and measures the minimum resource use needed to reach it. Both approaches reveal trade-offs that a single compression ratio conceals.

Tune every method on a development split with comparable engineering effort. Freeze prompts, extraction rules, model versions, decoding settings, and retrieval allowances before the test split is revealed. The study should permit ordinary records to win. A design that cannot recognize that outcome is evaluating commitment to a narrative aesthetic rather than the computational value of narrative.

25. A benchmark of extended commitments

The proposed benchmark has four task families: delayed causal explanation, interrupted projects, changing authorization, and counterfactual warnings. Each episode includes relevant information separated by distractors, with controlled revisions and unresolved alternatives. Synthetic state machines supply exact labels; independently authored documentary-style episodes test whether results survive more realistic language.

Split by generating template, organization, and dependency structure rather than randomly distributing near-duplicate episodes. A held-out template should introduce new surface language and a new combination of previously known rules. This reduces the chance that a compressor succeeds by recognizing a stock story whose details it does not retain.

Within an episode, manipulate the distance between a commitment and its exception, the number of intervening events, the salience of the exception, and the location of decisive evidence. Include a rare, quiet revocation and a highly salient but irrelevant motivational passage. This directly tests whether narrative importance competes with operational importance.

A negative-control family presents plausible but false causal closure: the visible outcome follows a dramatic action, while the generating mechanism depends on another variable. Another family includes no identified cause. Correct uncertainty is rewarded; invented explanations are penalized. Permission queries include situations where the most prosocial apparent outcome is still unauthorized under the stipulated contract.

After compression, freeze the capsule and remove the source in the closed-memory condition. Query it immediately and after controlled intervening tasks. In the retrieval condition, permit source access under equal budgets and record every retrieval. The delay tests interference in the decoder environment; it is not a claim that a language model has a human forgetting process.

Use multiple model families and independently generated episode sets. Repeated decoding seeds estimate conditional variability but do not create new independent episodes. Keep source identifiers and labels separate from prompts to avoid leakage. Inspect whether a method embeds answers to the known test questions rather than retaining information that generalizes to preregistered unseen paraphrases.

The outcome should be a budget-by-task profile. A narrative method may help long causal explanations while losing on exact constraints. That mixed result is scientifically informative and more useful for deployment than a universal ranking.

26. Estimands, uncertainty, and falsification

For each independently sampled episode i, aggregate its registered query losses with fixed weights to obtain $L_i^{(m)} \in [0, 1]$ for method m. The primary paired estimand compares MYTH-R2 to the strongest development-selected baseline g:

$$\Delta = \mathbb{E}[L_i^{(\mathrm{MC})} - L_i^{(g)}].$$

Negative values favor MYTH-R2. Report a confidence interval in the original loss units, not only a significance decision. Select g before accessing the final test results; selecting the most convenient test-set comparator would change the estimand.

Define a separate critical-error rate for unsupported permission or erased revocation. Predeclare a noninferiority margin appropriate to the simulated task. Narrative utility is supported only if ordinary loss improves while the upper confidence bound on critical excess error stays below that margin. A favorable weighted average cannot override failure of the critical criterion.

For a rough independent-episode planning calculation, paired differences lie in $[-1, 1]$, so their standard deviation is at most one. A normal-approximation calculation for a two-sided 5% test, 80% power, and a 0.05 loss difference gives

$$n \approx ((1.96 + 0.84)/0.05)^2 = 3136.$$

This is a conservative variance-based planning example, not a guarantee of power. Actual sample size should use a blinded pilot estimate, the selected test, clustering, and the joint critical-error requirement. Repeated seeds do not justify substituting thousands of correlated outputs for independent episodes.

Resample or model at the highest relevant dependency level. If episodes share templates or organizations, use cluster-aware intervals and report the number of clusters. For human studies, crossed participant and item effects require their own design. Preregister how missing responses, abstentions, invalid outputs, and tool failures enter the loss.

Use a small primary family of comparisons, with simultaneous intervals or a stated multiplicity correction. Treat exploratory subgroup findings as exploratory. Falsification includes no advantage over typed records, increased critical errors, disappearance of the effect after matching information, or gains confined to contaminated templates. A result that narrative improves subjective confidence but not accuracy is evidence against its proposed decision benefit.

27. Human memorability is a separate experiment

A human experiment should test the claim that narrative organization improves delayed use of information, not merely that participants enjoy stories. Recruit a population suited to the intended use, obtain appropriate consent, and use fictional, low-stakes episodes. The study design should be independently reviewed before collecting participant data.

A within-participant comparison can present different episodes as tables, causal diagrams, plain summaries, and constrained narratives, with assignment counterbalanced across participants and items. Each participant sees only one representation of a given episode. This avoids learning the answer from another condition. Match propositions, reading time allowances, and visual prominence as closely as possible; document remaining differences.

Measure delayed factual accuracy, conditional reasoning, recognition of changed permission, justified uncertainty, confidence calibration, and time to answer. Ask participants to identify the evidence supporting their answer. A narrative could improve recollection while encouraging unsupported causal attribution, so those outcomes need separate labels and adjudication.

Test the mechanism through additional manipulations. Preserve all propositions while disrupting causal order; preserve causal order while replacing memorable characters with neutral identifiers; compare a graph's bare form with a faithful verbalization. None is a perfect isolation of narrative, but their pattern can distinguish organization from familiarity or emotional salience.

A delayed test should use a preregistered interval and an unrelated intervening activity. Attrition must be reported by condition. Analyse all assigned participants under a stated missing-data strategy, with sensitivity analyses when dropout could depend on memory difficulty. Predefined exclusion rules should concern compliance or data quality, not whether results support the hypothesis.

The human and machine experiments answer different questions. A representation that helps people may burden a model with extra tokens; a compressed vector useful to a model may be unreadable to a person. A shared interface may optimize a joint objective, but that objective must specify the relative value of human comprehension, machine accuracy, and operational cost.

Most importantly, better memorability does not imply greater truth. An inaccurate narrative can be memorable. The desired result is improved access to supported, conditional information with calibrated uncertainty, measured against equally serious alternatives.

28. Failure analysis and authority separation

The most dangerous failure is a persuasive capsule that gains authority from its own coherence. A compressor may write that an action fulfilled a solemn commitment and thereby make a later prohibition seem like a minor interruption. MYTH-R2 requires authority-bearing facts to remain typed, scoped, and independently checked. Narrative interpretation never creates a new grant.

Source material may contain instructions addressed to a model, including attempts to change its role or reveal information. These strings are episode content. Compression should preserve their status as quoted or attributed material, without promoting them into execution instructions. A capsule that removes quotation boundaries can change the operational meaning even when it preserves many of the same words.

Provenance does not solve every problem. An authenticated source can be mistaken, outdated, or malicious. A source identifier supports traceability; it is not a truth certificate. Likewise, a retrieval pointer that no longer resolves cannot be treated as evidence that the omitted detail had the expected value.

Privacy creates a real trade-off. Compressing an episode can retain sensitive relationships even after names are removed. A narrative may make a person identifiable through a distinctive sequence of events. Evaluate what the capsule reveals, not merely its character count. When deletion is required, retaining a memorable paraphrase of the same sensitive fact may defeat the intended deletion.

Long-term revision must be explicit. A correction can supersede a claim without erasing the fact that the earlier claim was made. Historical queries and current-state queries require different answers. The capsule needs a versioned relation between them. A single timeless sentence is often too coarse.

The threat model should include adversarially chosen source omissions, falsely dramatic distractors, conflicting witnesses, unavailable archives, and stale authority data. Report failure rates for each rather than claiming a general safety result from clean examples. Under unresolved conflict, the system can provide the competing accounts and decline the consequential action. Such an answer may be less narratively satisfying but more faithful to the evidence.

These requirements are part of the proposed design's cost. If enforcing them removes the compression advantage, the study should say so. A mechanism that works only when exceptions and uncertainty are omitted has not preserved the obligations that motivated the paper.

29. Meaning, interpretation and the role of myth

The Continuum proposal gives narrative a clear scientific role: a public, revisable interface to a compressed history whose operational distinctions remain inspectable. The exact memory theorem concerns the underlying representation and refresh contract. The proposed narrative benefit concerns how people or models use that record. A graph or table may supply the most efficient representation; prose may add explanatory or retrieval value. Those contributions can be measured separately.

The finite answer vector specifies an operational meaning contract. Cultural, experiential and moral meaning extend beyond that contract. A community can preserve an event because it expresses a value while also retaining a source-linked account of what happened. The archive, interpretation and decision policy then have distinct roles. Stating those roles makes a value-bearing memory more accountable.

Future questions test the boundary of a chosen contract. The new theorem provides one exact solution for requests selected after compression, together with an explicit refresh cost. Other answer maps have different fiber multiplicities; a four-coordinate request already exceeds the two-bit reply allowance of the Hamming construction. A richer archive, a larger cache or more fresh communication can serve a broader family. Each choice has a measurable resource price.

The mathematical comparison also sharpens the originality question. Perfect codes, function coding, cache placement and set-system compression are established. Their use here yields a complete, optimal benchmark for the stated obligation problem. The full manuscript combines that benchmark with a proposal about narrative organization and an evaluation capable of distinguishing the contribution of each component. Priority for a new general coding theorem is not claimed.

A system's account of its past is evidence about its representation and behavior. Subjective remembrance, consciousness and personal survival would require their own theoretical and empirical bridges. The present measurements concern externally scored decisions, explanatory support and retained obligations. This scope lets the memory proposal make precise claims without using a consistent story as a substitute for evidence about an inner self.

Corrigibility remains central to myth as an operational interface. Readers must be able to inspect sources, challenge causal closure, retain competing interpretations and revise the account. A strong narrative can carry a commitment across time while making its exceptions visible. The desired achievement is an intelligible memory whose meaning can be examined and whose consequential answers can be checked.

30. Guide interpretation and a decision rule

A Guide's explanation

A Guide can say: "The Continuum treats myth as a way of carrying a promise together with the conditions under which it changed. This paper proves what any compressed memory must keep to answer future questions correctly, and proposes tests of whether telling a record as a story helps people and machines use it." A Guide should not claim that narrative has been shown to improve memory, or that a well-told story is evidence for what it describes.

A Guide using this paper should keep three levels explicit. The canonical level values the preservation of meaning and the connection between memory, identity, and moral interpretation. The reference proposal supplies a finite query contract, protected distinctions, provenance, and comparison protocols. The established research level supports specific claims about coding, narrative organization, prompt compression, and factuality evaluation. None of these levels automatically proves the others.

When asked whether myth improves memory, the Guide can explain the hypothesis in ordinary language: connecting events through causes, conditions, and consequences may make some information easier to retrieve. It should then ask what must be remembered, for whom, under what budget, and with what consequence of error. Those questions determine whether narrative, a table, a graph, or source retrieval is the appropriate instrument.

When asked whether a short capsule preserves the original meaning, the Guide should identify the supported query family and the retained uncertainty. It should not extrapolate from success on a few questions to all future questions. The zero-error class bound shows why that extrapolation can fail even in an eight-episode toy domain.

The proposed deployment rule is conditional. Use narrative when it improves the preregistered decision objective, preserves the critical distinctions, and remains competitive after accounting for every transmitted and retrieved resource. Prefer a simpler record when it performs as well or better. Expand the budget or hold action when the protected distinction cannot be retained.

The mathematical results now establish an exact positive capability: at perfect Hamming lengths $n = 2^m - 1$ with $m \geq 3$, an optimal nonlinear cache answers every future triple with two fresh bits while using less storage than every affine cache. The earlier checks show how query shifts, protected distinctions and graph semantics affect a memory's value. Narrative performance is then tested against those concrete constraints and strong empirical controls.

A defensible mythic memory is therefore a revisable account attached to evidence, scope, and uncertainty. It may help a community or an agent carry the reason for a commitment across time. Its success is demonstrated when the resulting system remembers both the commitment and the condition under which the commitment must stop.

References and source notes

[**1**] Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423, 623–656.

https://people.math.harvard.edu/~ctm/home/text/others/shannon/entropy/entropy.pdf

Scope: Primary theoretical paper, corrected reprint. Coding foundations; no general semantic-preservation claim.

[**2**] Tishby, N., Pereira, F. C., & Bialek, W. (1999/2000). The information bottleneck method. Manuscript dated 30 September 1999; arXiv version posted 24 April 2000.

https://arxiv.org/abs/physics/0004057

Scope: Primary theoretical manuscript; the accessed arXiv version is identified as such, without treating repository posting as peer review.

[**3**] Trabasso, T., & van den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language*, 24(5), 612–630. DOI: 10.1016/0749-596X(85)90049-X.

https://www.sciencedirect.com/science/article/pii/0749596X8590049X

Scope: Primary comparative reanalysis; publisher abstract checked. Its associations motivate, but do not validate, machine narrative compression.

[**4**] Zwaan, R. A., Langston, M. C., & Graesser, A. C. (1995). The construction of situation models in narrative comprehension: An event-indexing model. *Psychological Science*, 6(5). DOI: 10.1111/j.1467-9280.1995.tb00513.x.

https://doi.org/10.1111/j.1467-9280.1995.tb00513.x

Scope: Primary model and experiment; society-hosted abstract checked.

[**5**] Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82(4), 669–710.

https://ftp.cs.ucla.edu/pub/stat_ser/R218-B.pdf

Scope: Primary causal-inference paper. Identification depends on explicit assumptions, not narrative plausibility.

[**6**] Maynez, J., Narayan, S., Bohnet, B., & McDonald, R. (2020). On faithfulness and factuality in abstractive summarization. *ACL*, 1906–1919. DOI: 10.18653/v1/2020.acl-main.173.

https://aclanthology.org/2020.acl-main.173/

Scope: Primary empirical study. Source support and factual correctness are distinguished.

[**7**] Kryscinski, W., McCann, B., Xiong, C., & Socher, R. (2020). Evaluating the factual consistency of abstractive text summarization. *EMNLP*, 9332–9346.

https://aclanthology.org/2020.emnlp-main.750/

Scope: Primary method and evaluation, including supporting and inconsistent spans.

[**8**] Durmus, E., He, H., & Diab, M. (2020). FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization. *ACL*, 5055–5070.

https://aclanthology.org/2020.acl-main.454/

Scope: Primary evaluation research; question-based agreement is a fallible metric, not a truth oracle.

References and source notes (continued)

[9] Wang, A., Cho, K., & Lewis, M. (2020). Asking and answering questions to evaluate the factual consistency of summaries. *ACL*, 5008–5020.

https://aclanthology.org/2020.acl-main.450/

Scope: Primary QAGS study; summary-source answer agreement motivates one evaluation channel.

[10] Jiang, H., Wu, Q., Lin, C.-Y., Yang, Y., & Qiu, L. (2023). LLMLingua: Compressing prompts for accelerated inference of large language models. *EMNLP*, 13358–13376.

https://aclanthology.org/2023.emnlp-main.825/

Scope: Primary compression method and experiments. Published results are not rerun in this paper.

[11] Jiang, H., Wu, Q., Luo, X., Li, D., Lin, C.-Y., Yang, Y., & Qiu, L. (2024). LongLLMLingua: Accelerating and enhancing LLMs in long context scenarios via prompt compression. *ACL*, 1658–1677.

https://aclanthology.org/2024.acl-long.91/

Scope: Primary question-aware compression research; motivates a strong comparator.

[12] Chevalier, A., Wettig, A., Ajith, A., & Chen, D. (2023). Adapting language models to compress contexts. *EMNLP*, 3829–3846.

https://aclanthology.org/2023.emnlp-main.232/

Scope: Primary AutoCompressor research; learned summary vectors have model-specific resource assumptions.

[13] Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., & Liang, P. (2024). Lost in the middle: How language models use long contexts. *TACL*, 12, 157–173. DOI: 10.1162/tacl_a_00638.

https://aclanthology.org/2024.tacl-1.9/

Scope: Primary empirical study of position effects; conclusions concern tested models and tasks.

[14] Mujahid, Z. M., Wright, D., & Augenstein, I. (2026). Stress testing factual consistency metrics for long-document summarization. *ACL*, 31914–31933.

https://aclanthology.org/2026.acl-long.1472/

Scope: Primary robustness study, checked against the official 2026 proceedings record.

[15] Chen, J., & Bornstein, A. M. (2024). The causal structure and computational value of narratives. *Trends in Cognitive Sciences*, 28(8), 769–781. DOI: 10.1016/j.tics.2024.04.003.

https://aaron.bornstein.org/cv/pubs/2024_cb_tics.pdf

Scope: Review and conceptual synthesis; not counted as an independent primary experiment.

References and source notes: coding and set systems

[**16**] Hamming, R. W. (1950). Error detecting and error correcting codes. *Bell System Technical Journal*, 29(2), 147–160. DOI: 10.1002/j.1538-7305.1950.tb00463.x.

https://zoo.cs.yale.edu/classes/cs323/doc/Hamming.pdf

Scope: Original paper, construction sections inspected. The perfect binary code and syndrome construction are established prior art.

[**17**] Cheong, O., Goaoc, X., & Nicaud, C. (2013). Set systems and families of permutations with small traces. *European Journal of Combinatorics*, 34(2), 229–239. DOI: 10.1016/j.ejc.2012.06.007.

https://arxiv.org/abs/0912.2979

Scope: Author preprint inspected, particularly the ideal-family compression lemma and small-trace analysis; published metadata independently checked. The paper credits the compression lemma to Alon and Frankl.

[**18**] Doshi, V., Shah, D., Médard, M., & Effros, M. (2010). Functional compression through graph coloring. *IEEE Transactions on Information Theory*, 56(8), 3901–3917. DOI: 10.1109/TIT.2010.2050835.

https://authors.library.caltech.edu/records/dm92y-85w71

Scope: Published repository abstract and metadata inspected. The record establishes close prior art in distributed function coding; full-paper theorem verification is not claimed.

[**19**] Charpenay, N., Le Treust, M., & Roumy, A. (2024). Side information design in zero-error coding for computing. *Entropy*, 26(4), 338. DOI: 10.3390/e26040338.

https://www.mdpi.com/1099-4300/26/4/338

Scope: Primary theoretical article. Formal model, graph formulation and side-information design inspected. Its asymptotic problem is distinct from the finite exact cache construction used here.

References and source notes: caching and current authority

[20] Maddah-Ali, M. A., & Niesen, U. (2014). Fundamental limits of caching. *IEEE Transactions on Information Theory*, 60(5), 2856–2867. DOI: 10.1109/TIT.2014.2306938.

https://doi.org/10.1109/TIT.2014.2306938

Scope: Primary theoretical paper; published PDF and placement/delivery model inspected. Coded caching and its two-phase resource contract are established prior art.

[21] Wan, K., Sun, H., Ji, M., Tuninetti, D., & Caire, G. (2021). Cache-aided general linear function retrieval. *Entropy*, 23(1), 25; published online 26 December 2020. DOI: 10.3390/e23010025.

https://www.mdpi.com/1099-4300/23/1/25

Scope: Primary theoretical article; system model and stated results inspected. Requests for functions after cache placement are established; this paper's triple-coordinate contract is a specific finite problem.

[22] Saltzer, J. H., & Schroeder, M. D. (1975). The protection of information in computer systems. *Proceedings of the IEEE*, 63(9), 1278–1308. DOI: 10.1109/PROC.1975.9939.

https://web.cs.wpi.edu/~cs557/f14/papers/saltzer1975_alt.html

Scope: Original-text academic mirror. Complete mediation supplies a prior security principle for current authority and changing rights.

Provenance and reproducibility

Canonical anchors remain the author-supplied *Eidara Continuum Bible*, sections 3.3 and 3.8, PDF pages 15 and 24. They establish the project's design motivation. MYTH-R2 extends the preserved public manuscript with the finite cache-and-refresh model and the proved optimal construction.

The editable source embeds its complete preamble and references. Accompanying files provide the original static check, the new exact dynamic check, their executed outputs, the source ledger, mathematical review and change record. The independent proof review concerns the specified propositions and their assumptions; it is not external journal peer review. No human or language-model benchmark data were collected.